

Planning to Explore: Using a Coordinated Multisource Infrastructure to Overcome Present and Future Space Flight Planning Challenges

Edward Balaban^{*}

Michael Orosz^{**}

Tatiana Kichkaylo^{**}

Andre Goforth^{***}

Adam Sweet^{*}

Robert Neches^{**}

^{*} QSS Group Inc./NASA Ames Research Center, Moffett Field, CA 94035

^{**} University Of Southern California Information Sciences Institute
4676 Admiralty Way, Marina Del Rey, CA 90292

^{***} NASA Ames Research Center, Moffett Field, CA 94035

Abstract

Few human endeavors present as much of a planning and scheduling challenge as space flight, particularly manned space flight. Just on the operational side of it, efforts of thousands of people across hundreds of organizations need to be coordinated. Numerous tasks of varying complexity and nature, from scientific to construction, need to be accomplished within limited mission time frames. Resources need to be carefully managed and contingencies worked out, often on a very short notice.

From the beginning of the NASA space program, planning has been done by large teams of domain experts working months, sometimes years, to put together a single mission. This approach, while proven very reliable up to now, is becoming increasingly harder to sustain. Elevated levels of NASA space activities, from deployment of the new Crew Exploration Vehicle (CEV) and completion of the International Space Station (ISS), to the planned lunar missions and permanent lunar bases, will put an even greater strain on this largely manual process. While several attempts to automate it have been made in the past, none have fully succeeded. In this paper we describe the current NASA planning methods, outline their advantages and disadvantages, discuss the planning challenges of upcoming missions and propose a distributed planning/scheduling framework (CMMD) aimed at unifying and optimizing the planning effort. CMMD will not attempt to make the process completely automated, but rather serve in a decision support capacity for human managers and planners. It will help manage information gathering, creation of partial and consolidated schedules, inter-team negotiations, contingencies investigation, and rapid re-planning when the situation demands it. The first area of CMMD application will be planning for Extravehicular Activities (EVA) and associated logistics. Other potential applications, not only in the space flight domain, and future research efforts will be discussed as well.

Introduction

The ability to generate and maintain plans and schedules in a multi-discipline and highly distributed environment is a subject of much research [1]. Examples of such environments include military and civilian aviation, commercial trucking and freight operations, and manned/unmanned space flight operations. Typically, stakeholders in these environments collaborate and generate plans and schedules using a collection of disparate tools and manual techniques. Although these methods are sufficient in many environments, they do not scale well as the size and complexity of the problem increases. Typically, as the size and complexity of the domain grows, the quality of the generated plans and schedules suffer. This is particularly true in manned space flight environments where plans and schedules undergo constant modifications due to changing operating conditions. Under these operational conditions, traditional methods suffer from many problems including: scheduling conflict resolution, communication breakdowns, lack of situational awareness, difficulty in undertaking schedule repair as opposed to re-plan, and lack of support for exploring alternative planning and scheduling solutions.

Coordinated Multi-source Maintenance on Demand (CMMD) is a multi-agent planning and scheduling decision support tool currently under development for future NASA manned and robotic space flight operations. CMMD is not an autonomous planning/scheduling tool. Rather, CMMD is a planning, scheduling, and execution decision-support system that aids NASA planners in developing and maintaining plans and schedules.

A key feature of CMMD is the integration of maintenance into the overall planning and scheduling process. Currently, NASA addresses this problem through manual

negotiations among disciplines. This is sufficient as long as the size of the problem is small. Unfortunately, as NASA takes on more planning as part of the new Space Exploration Program, the complexity of these negotiations will greatly increase. This can lead to planned/scheduled events being cancelled or aborted due to a variety of factors including lack of resources, resource over-use, and incorrect resource allocation. In CMMD, planners and schedulers from all areas collaborate through negotiation technologies to achieve a coordinated plan/schedule that meets the needs of the mission.

Other key features of the system include generating so-called “good enough, soon enough” plans, legal schedules, localized planning and scheduling, and the ability to undertake multiple “what-if” analysis. CMMD is currently being developed to support NASA’s Space Exploration Program where multiple disciplines distributed over multiple localities must collaborate in the planning and scheduling of science/exploration, habitat build-out, and maintenance tasks and operations. Although targeted for the Space Exploration Program, the initial prototypes will, as a proof-of-concept, be used to support increment planning and scheduling of ISS operations.

In the following sections the NASA domain, the current concept of operations, and the planning challenges will be described, followed by a technical description of CMMD and its proposed role in future NASA space activities.

Current Planning Approach for Space Missions

Domain Background

The application domain which this effort is centered around is NASA’s manned space operations led by the Johnson Space Center (JSC) in Houston, Texas. The two major programs, currently well underway, are the Space Shuttle and the ISS. The overall mission for these programs has been to provide manned presence in low earth orbit for the conduct of space operations such as satellite servicing and scientific research. The third major program is Exploration Systems, now in the formulation stage, that has the mission of returning humans to the moon and ultimately taking them to Mars. It will eventually subsume the existing programs, where the Space Shuttle will be replaced with the CEV and the ISS will serve as a test-bed for validating systems and technologies for crewed lunar and planetary outposts, including the technologies developed in this effort.

While the Space Shuttle program has long established operational processes, procedures and planning/scheduling tools, the ISS is still in the process of construction and so its processes have to accommodate the unique, one time construction activities for completion of station assembly. As a consequence, its operational processes must accommodate the growth of the ISS, as new components

are added and additional participation of international partners is accommodated.

The organizational structure at JSC responsible for these programs has its roots in the days of the Apollo space program and, because of this successful legacy and proven track record, will likely be adapted, with incremental modifications, to meet the needs of Explorations Systems.

This organization, with approximately 3000 people involved, has flying missions as its primary objective. The division of labor to support this objective is divided into flight operations, planning, training and facilities. The major functions within the flight operations area are real-time operations (involving flight directors and controllers), integration of operations with international partners, and real-time planning and scheduling.

The planning activities consist of mission concept definition, mission requirements integration, long-term and short-term planning of Systems, EVA and Intra-Vehicular Activity (IVA) tasks, robotics activities, payload and launch integration, and several others.

CMMD work is initially focused on ISS operations and logistics planning, with a particular focus on EVA tasks; however, support for the execution area and its real-time specifics will be incorporated in later phases. The section below describes in detail how operations planning is done currently in order to provide the reader with the understanding of motivation guiding CMMD work.

ISS Operations Planning

In planning ISS operations, the largest unit of time ordinarily dealt with is an *increment*. An increment is the period of time covering activities of a specific crew onboard a space station or at a habitat. The currently used duration of ISS increments is roughly six months.

Initial Increment Definition Phase. Planning for an increment starts with the assignment of an Increment Manager (IM). This person is the final authority on everything having to do with increment planning and also plays a significant role when execution of the increment begins.

The IM assembles the Increment Management Team (IMT) consisting of representatives from various disciplines involved in the mission, such as engineering, science, medical, EVA, and others. This happens a year to a year and a half before the actual increment starts. The group’s first objective is to come up with the Increment Definition Requirements Document (IDRD). This document defines, in broad strokes, the major stages of the increment, such as manned and cargo spacecraft arrivals and departures, as well as primary science, assembly, and maintenance tasks. Some event dates are defined but assumed to be approximate, unless related to a time-sensitive event such as a launch.

IMT is guided in the IDRD creation process by the following: the increment definition statement (usually a paragraph or two defining the main goals of an increment),

a document describing the strategic plan for the ISS, the Generic Ground Rules & Constraints (GGR&C) document defining the common rules for ISS activities, the ISS assembly sequence plan, and the documents of the Crew Time Working Group. The later specifies the approximate distribution of crew time among the major types of activities (such as research, maintenance, education, training, and others). Crew time allotments, arguably being the most valuable resource on the ISS, are often a subject of intense negotiations among the discipline teams. Finally, IMT maintains a list of proposed tasks, submitted by stakeholders within NASA and by the outside scientific community. This list is prioritized and new tasks go through a formal selection process conducted by the IMT.

On-Orbit Summary (OOS). The next level of planning details is defined in the OOS document. The work on OOS starts about six to four months before the actual start of the increment and often overlaps in time with the final stages of IDRD completion. OOS is created by the planners from the Mission Operations Directorate (MOD) of JSC with the assistance of Subject Matter Experts (SME) from other groups, who advise the planners on engineering, environmental, medical, and other issues.

OOS covers the duration of the entire increment and describes it on the day-by-day basis. Tasks for each day are listed, although the order and the duration of tasks is kept flexible. There is also a practice of inserting “padding” between tasks to help deal with variations in crew working styles as well as with the situations where a task takes much longer than expected.

The main inputs into OOS are usually the IDRD, the GGR&C, and also the documents describing discipline-specific rules and constraints.

Weekly Look-ahead Plan (WLP). Once an increment begins, the Increment Management Team (IMT) transitions from planning duties to overseeing the execution of the increment. It also changes its name to Increment Management Center (IMC). While the IMC is not actively involved in further planning activities – they are performed by MOD planning teams – it gets called upon when major tasks need to be added, cancelled, or rescheduled.

Once the execution of an increment begins, the planning teams take a week-long section of OOS two weeks before it starts and plan it out on the next level of detail.

Tasks within each day are ordered and get time slots allocated for them. Personnel and resources get assigned to specific tasks and in some cases complex tasks are broken up into subtasks for scheduling purposes. This becomes the Weekly Look-ahead Plan (WLP). Negotiations among the various teams and international partner representatives trying to get their tasks accomplished and constraints satisfied take place at the WLP meetings. These negotiations tend to get particularly heated towards the end of an increment, when the opportunities to get tasks accomplished start to become

scarce, with the crews preoccupied with preparing for handover and return to Earth.

Short Term Plan. The final level of detail is provided in the Short Term Plan (STP). The planners produce it by taking a day out of the WLP a week in advance and solidifying the times/ordering for tasks, crew assignments, and other details. The plan is then circulated among the international partners and stakeholders for approval. The approval process takes roughly a day.

Next, the plan is recorded in an electronic form and becomes the On-Orbit Short Term Plan (OSTP). It is then uploaded onto the on-board computers. Once execution of an OSTP begins, the crew can mark the completion status of each task and provide comments, such as comparing the expected level of difficulty to the actual. Status information is also collected by the ground personnel at the end of each day via a videoconference and is used to adjusting the next day’s OSTP, if needed.

Tools Utilized. Currently the planning process described above is very laborious and manual, especially when unexpected changes need to be integrated. Attempts to automate it have been made, however, the systems deployed proved to be cumbersome to use and lacking in features. The currently utilized Consolidated Planning System is primarily being used only as a schedule recording, viewing, and distribution tool. The tools used by the planners most often are actually Microsoft Excel spreadsheets and Word documents. Most of the supporting information for the planning process is stored either in such spreadsheets and documents, or in largely stand-alone databases, maintained by various groups. This situation obviously creates difficulties with data synchronization and knowledge transfer. This, and other challenges facing mission planners are covered in the next section. The CMMD section will then discuss how these challenges are being addressed in our system.

Challenges

Challenges facing manned spaceflight

In the next 5-8 years NASA will face the challenge of simultaneously undertaking three major programs. First, NASA will be completing the build-out of the ISS while simultaneously performing various ISS-related maintenance and science operations. Secondly, NASA will continue to operate the Space Shuttle to finish the ISS construction and to possibly to fly specialized missions, such as Hubble Space Telescope repair or dedicated science missions. Finally, NASA is developing, testing, and deploying two new types of spacecraft: Crew Exploration Vehicle (CEV) and Crew Launch Vehicle (CLV). The CEV will first be deployed as an ISS lifeboat, then used as an integral part of the NASA’s new Exploration Program for missions to the lunar surface and, eventually, to Mars. Below is an overview of the challenges that will arise in the near and longer terms for

space exploration and their implications for planning/scheduling work:

Increased level of activities. In the near-term, ISS build-out activities will put a severe strain on the current planning/scheduling infrastructure[2]. For instance, instead of the typical 4-5 scheduled EVAs per year, it is anticipated that between 25-30 will be required. In addition, as the build-out occurs, equipment (such as batteries or gyroscopes) already installed and operational will, in many cases, approach its operational life-time limits. This will, of course, result in more maintenance EVAs being scheduled in addition to the increased number of construction EVAs.

In addition to the ISS build-out and on-going maintenance needs, new robotic equipment will be deployed on the station. This equipment will augment EVA and IVA tasks and, therefore, offer another level of planning complexity. For example, planning and scheduling will now include decisions over whether to schedule a human or a robot for a particular EVA. Issues such as increased robotic maintenance will also add to planning/scheduling complexities.

As new modules are added during the ISS build-out, such as Columbus science module from the European Space Agency (ESA), there will be increased scheduling of scientific tasks. Crew size will also most likely increase to accommodate the new science missions and to support the increased ISS maintenance activities. While it is unlikely that the station's crew size will increase to the originally planned six to seven-person crews, it is anticipated that the current two-person crew will be expanded to a four-person crew to allow greater participation by international partners.

New types of spacecraft. Several new spacecraft types are nearing deployment and will have to be integrated into the planning and scheduling process. These new spacecraft include ESA's Automated Transfer Vehicle (ATV), the aforementioned CEV, and possibly Russia's new Kliper – a crew-carrying spacecraft that will also be used as a station lifeboat. In each case, the planning and scheduling process must take into account a variety of new information, such as payload capacity, size of crew, frequency of launch, length of flight, and other specialized needs and capabilities.

Logistical difficulties. With the availability of the US Space Shuttle fleet to undertake frequent flights in question, sustaining an effective logistical pipeline to and from the station will be difficult. Currently, the other available cargo spacecraft, the Russian Progress, has a limited capacity (2.7 metric tons vs. Shuttle's 25 tons) and cannot return cargo from orbit. The ESA's ATV, while increasing the amount of cargo that can be launched to approximately 9 metric tons, will, like the Russian Progress, not be able to return cargo from orbit. The CEV and the Kliper will be capable of returning back some cargo, however, it will most likely be on the order of a few hundred kilograms and will be limited in volume. All of

these factors will need to be taken into account when planning logistical support for larger crews, an aggressive ISS build-out schedule, and an increased science agenda.

The Moon-Mars Exploration Program, in addition to dealing with many of the present challenges, will also present some new ones:

Planetary base operation. While NASA and its international partners have by now accumulated significant experience operating space stations in low Earth orbit, no manned planetary bases have ever been established. NASA plans to operate such bases on the lunar surface, where they would be manned for extended periods of time, and use this experience as the basis for eventual expeditions to the Martian surface. It is assumed that operating a lunar base will be similar in many respects to operating a space station, but there will also be significant differences, such as those related to in-situ resource utilization and lengthy surface excursions, possibly with the help of robotic assets.

Greater operational distances. While multi-million mile expeditions to Mars are at least a couple of decades away, even lunar operations, conducted at a distance of a "mere" 240000 miles, will present significant obstacles in terms of logistical support, communications, and situational awareness of the ground support personnel. To accommodate these challenges, more of the planning and scheduling authority (particularly in the day-to-day operations) will need to be given to astronaut crews. This will increase the complexity of the overall planning and scheduling process since both local and remote planning/scheduling needs and requirements will have to be integrated – a particularly challenging prospect given the distance and anticipated communication breakdowns between ground and lunar operations.

Operating multiple simultaneous missions. Another new challenge will be the planning of simultaneous, possibly interdependent, multiple missions. Examples include a crew operating in lunar orbit and another crew in a lunar habitat, or possibly crews operating from two separate habitats. This will raise the issues of task coordination and resource sharing, and will require rethinking of the current planning/scheduling approach as well.

Challenges in adopting automated planning systems in space flight domain

During the collection of domain information for the CMMD project, several challenges were discovered which negatively impact the complete adoption of any automated system for planning spaceflight activities:

"Living in a house while it is being built". In the case of the ISS, and likely in the case of future complex spacecraft and habitats, the vehicle will still be undergoing build-out while simultaneously being used. In an interview at JSC, a member of the ISS activity planning group stated "the process has changed since we began, and it will probably

change again before we are finished". Thus, any automated planning system will have to be flexible and adaptable. The priority of objectives and activities will also change over time and this must be accommodated as well.

Planning process is very manual. The current planning process is based on manual efforts of many people at JSC and throughout NASA. Only a small portion of the knowledge required to accomplish the planning task is in a format suitable for machine reasoning. The majority of the knowledge base exists in documents (both paper and electronic) and, in many cases, only in the memory of the planning personnel. Even the planning process itself is not documented in great detail, in part because it is changing over time. Many of the vehicles and equipment in the space program are one-of-a-kind, and only have a few people knowledgeable in their use and maintenance. Finally, the human planners have expertise and intuition gained from many years of planning and executing ISS operations. A large body of knowledge incorporating the above factors must be codified into a form suitable for machine reasoning by an automated planning system.

Large, complex operating space. The distributed nature of the planning process, the special knowledge of the planners, the sheer size of the planning/scheduling problem, and number of personnel involved will be difficult to duplicate in a fully automated planning system. This complexity will only increase as the number of habitats, transport vehicles, and science missions increase. As noted earlier, the planning/scheduling system must scale with the size of the problem.

Need for human involvement. Human personnel, at least for the foreseeable future, are ultimately responsible for the increment plan/schedule and the health and safety of the crew. If an automated planning/scheduling system is to be successfully deployed, it must produce plans and schedules that are correct and safe to implement. This, of course, will require human involvement and, in many cases, user-directed plan/schedule adjustments. Therefore, it is likely that any planning system will have to be mixed-initiative, taking at least some guidance from humans. Also, an automated system should be able to help humans analyze alternatives and "what-if" scenarios, leaving ultimate decisions up to them. Traceability and an explanation facility will be key features to show operators the feasibility of particular plans and schedules. Finally, a system able to repair existing plans is highly desirable; it will allow long, "living" plans to exist, modified as needed over time.

Several previous attempts to automate the planning process have failed because they were over-precise, did not involve human planners throughout the entire planning cycle, nor were the planning/scheduling algorithms sufficiently mature to handle complex planning problems. Based upon these experiences, the following requirements were formulated for the proposed CMMD system:

- The system must be distributed across the enterprise to allow participation of multiple groups in various areas of expertise
- The system needs to be able to incorporate a variety of information into the planning process, from flight rules and constraints, to technical information, to human preferences and heuristics
- It must allow human interaction throughout the planning process, while relieving planners from tedious manual tasks, such as entering data or doing manual re-planning
- Related to above, it needs to have a fast, localized plan repair capability to minimize the impact of problems and delays in one area on the rest of the schedule
- The system must be adaptable and allow rapid input of new information and one-time events, human preferences and configuration changes
- In assisting with planning/scheduling tasks, it needs to address the associated logistics issues
- The system must be able to deal with communication interruptions from ground information sources when deployed on-board a manned spacecraft
- The system needs to minimize time and effort that is currently spent on inter-group negotiations
- Finally, it needs to provide options to investigate (i.e., "what if" capability)

CMMD

The goal of the Coordinated Multi-source Maintenance on Demand (CMMD) project is to build a system that addresses the challenges described in the previous section.

A CMMD system consists of multiple independent agents connected to a virtual Backbone. The Backbone maintains a representation of the current state of the plan/schedule, called the Living Schedule (LS). In addition, the Backbone also contains the current state of various cargo manifests and habitat inventories. The Backbone performs access control and notification propagation allowing agents with different areas of expertise to collaborate on the LS, manifests, and inventories.

Domain information, such as rules, constraints, preferences, and variables is collected through the PRISM (Planning & Real-time Information for Space Missions) component that connects to already existing databases (such as JSC's Orbital Data Reduction Complex) and supplements it with user entered data. The information is then processed, organized, and indexed in an internal domain data repository. As needed, it is transmitted to the agents (and LS) via Backbone and translated into the appropriate representation.

Information is contained in the LS in a declarative form. The CMMD data model separates representation of the plan/schedule from rules and procedures used to modify the plan and also from user preferences. In addition, the LS stores dependencies between elements of the plan, so that

every planning decision made by CMMD is traceable to the rule or preference this decision is based on. Dependencies between elements of the plan and the contents of cargo manifests and habitat inventories are also maintained.

Such cause-effect dependencies can be individually inspected and disabled by users. CMMD propagates the effects of any changes made by the user. This way the user has complete information and control over planning decisions, while the CMMD system handles the tedious and error-prone dependency propagation.

Note that dependency propagation performed by CMMD takes into account logistics rules and constraints. This way, when a user adds a new task to the schedule, the system automatically computes logistics implications of this task and assesses its feasibility, cost, and impact on other activities.

The size of the plan stored in the LS and the total number of agents working on it may be quite large. Finding an optimal solution would require a complete search, which is infeasible in such a large problem. Moreover, the problem with undertaking a complete search of the solution space is compounded by constantly changing conditions and user preferences. CMMD addresses this issue through algorithms that focus on finding “good enough” solutions, plan repair, and evaluation of alternatives.

Finally, CMMD is a multi-agent system. The virtual Backbone that connects the agents is also responsible for storing the Living Schedule and related cargo manifests and habitat inventories, as well as for notifying agents about relevant changes in the LS. Various services provided by the Backbone allow agents to cope with temporary disconnections and facilitate negotiations by keeping track of dependencies while minimizing traffic.

In the rest of this section we discuss various elements of the CMMD architecture in more detail.

Domain and plan representation

CMMD relies on a data model that combines features of constraint-based systems ([3], [4], [5]) with those of hierarchical task network-based (HTN) planners ([6], [7], [8]). In CMMD, physical and virtual resources are described in terms of timelines. Activities that use these resources are represented by tokens placed on the timelines. Relationships between the activities are captured by procedural constraints connecting variables of the tokens. For example, one may restrict the end time of one activity to be less than the start time of another activity (the temporal *Before* constraint).

In addition to resource timelines, the CMMD data model supports real-valued capability timelines. Capabilities are tied to resources. For example, resource Battery may have capability Charge, whose value is affected by Use and Recharge tokens placed on the Battery timeline.

At any given moment, the state of the CMMD plan is described by the tokens residing on the various resource

timelines. A user can change the plan by explicitly adding and removing tokens. Often such a change brings the plan to an invalid state, which can be fixed by adding, removing, and/or modifying other tokens. For example, scheduling a particular crew member to perform an EVA at some point in the future requires scheduling the same crew member for training before the proposed mission.

One of the goals of our system is to help the user keep track of such dependencies between activities. CMMD captures such dependencies using two kinds of rules: safety and achievement rules. Safety rules describe legal states of the world, while achievement rules represent standard expansions of high-level activities into procedures.

As with most constraint-based planning systems, CMMD allows non-singleton values for variables at different stages of the problem solving process. A distinguishing feature of the CMMD representation is that *all* choice points are captured using variables. This includes, among other things, the choice of a resource to be used in a particular activity (*timeline* variable of a token) and the choice of expansion rule for a high-level activity (*support* variable of a token).

Multi-agent architecture

A CMMD system consists of multiple agents connected to the Backbone. Backbone is a virtually centralized entity. However, it can be implemented using either a centralized solution or peer-to-peer technology.

The Backbone provides the following functionality:

- Persistence and 24x7 access to the Living Schedule, cargo manifests, and habitat inventories
- Session management and access control for agents
- Query and subscription service

For the most part, agents represent the various planning disciplines in the current NASA planning process. However, neither the set, nor the types of agents connected to the Backbone are fixed. In a typical mission planning environment, each human user would have a dedicated agent. Legacy systems and external data sources would also have their own agents. Various physical resources might also be connected to the Backbone using agents.

All CMMD agents implement the same interface to connect to the Backbone. Agents can read and write different portions of the Living Schedule as determined by their access rights. For example, an agent corresponding to a physical resource submits real-time data about the state of this resource to the Living Schedule. Agents interested in the state of this resource post a subscription on that resource to the Living Schedule. Upon receipt of an update from the physical resource agent, the Backbone checks the list of current subscriptions and sends notifications to all interested agents.

In addition to querying and changing the state of the Living Schedule, agents can negotiate with each other by

exchanging proposed “patches” to the current state of the Living Schedule before committing their changes. This process is discussed in more detail below.

CMMD agent

CMMD agents can operate independently or as part of a multi-agent community by communicating and negotiating with each other via the Backbone.

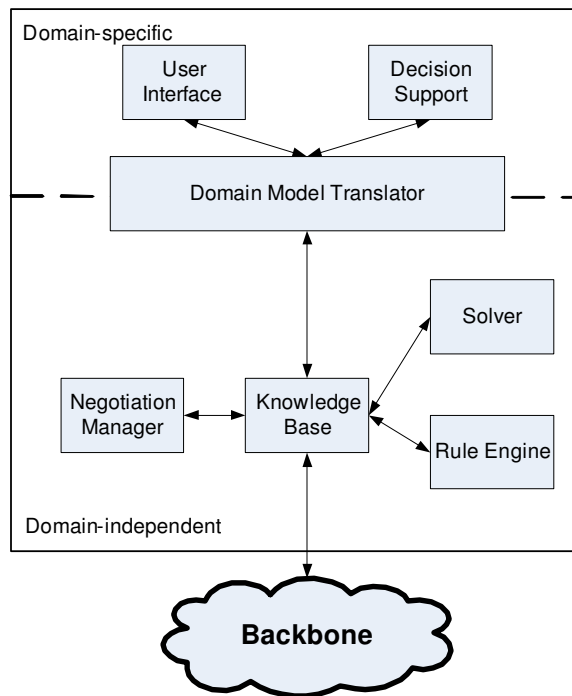


Figure 1. CMMD Agent Architecture

The figure above shows a simplified structure of a CMMD agent. Agent modules can be divided into two classes: domain-independent and domain-specific.

At the core of the domain-independent portion of the agent sits the Knowledge Base (KB). The KB stores the local view of the data from the Living Schedule, the agent’s internal data, and intermediate results of computations. The KB is also responsible for communicating with the Backbone and performs other functions, such as detecting inconsistencies and keeping track of dependencies among alternative solutions.

The Rule Engine (RE) relies on safety and achievement rules to help ensure the validity of the plan. When modifications are made to the plan, the RE checks the state of the plan stored in the KB and fires applicable safety and achievement rules. In many cases, these rule firings may cause the plan to be updated to ensure the validity and legality of the plan.

The Rule Engine is a purely production system. It does not perform search or any expensive computations. The Solver, on the other hand, is responsible for finding a consistent assignment of values to all variables in the plan – this requires search. CMMD architecture does not

prescribe which algorithm to use for this purpose. Currently we are investigating several implementations of the Solver module, including algorithms that accept external guidance in the form of variable and value ordering heuristics.

As mentioned earlier, CMMD agents can evaluate multiple alternatives. Moreover, agents can “discuss” alternatives with each other before submitting them to the Living Schedule. For example, during the planning of an EVA, the EVA agent may consult with the Maintenance agent and the Flight Surgeon to confirm that both the equipment (e.g., space suit) and the crew members are ready for the space walk. The CMMD system allows for “discussing” multiple alternatives with different agents without involving human users. These localized negotiations (undertaken by the Negotiation Manager) eventually generate and present different options – some more feasible than others – to the user. The benefit of this negotiation process is that the agents do most of the work, leaving the human decision-maker to make the final choices.

All modules described so far are domain-independent. The same code can be used in other domains in addition to space exploration. The only difference will be the information stored in the Living Schedule and the sets of rules used by the agents. Some modules, however, are necessarily domain-dependent. The most obvious example is the user interface. A CMMD agent can also include a Decision Support module responsible for generating suggestions and detecting opportunities based on the current state of the schedule.

The domain-specific modules of the architecture read and modify information stored in the Knowledge Base via the Domain Model Translator (DMT). The DMT permits more natural access to domain objects without jeopardizing reusability of the domain-independent modules.

CMMD Role in Future NASA Planning Operations

As noted earlier, the current NASA planning organization will be gradually adapted to meet the needs of NASA’s up-and-coming Exploration Systems program. CMMD’s approach to being a part of this shift is to serve as a virtual organization that overlays the existing one. While there are a number of decision support and analysis tools already in use to assist in the planning and scheduling of operations, their flexibility in incorporating a wide variety of information sources and adapting to constant change is limited.

The impact of CMMD on future operations is envisioned as being a provider of a revolutionary level of flexibility within the existing model of operations. Users will be able to conveniently incorporate pertinent information into the planning process, conduct comprehensive inter-discipline negotiations, and rapidly evaluate a large number of “what-if” situations, when needed.

One of the major challenges in carrying out space exploration is maintaining a reliable logistics pipeline. The current logistics concept of operations is adequate for Low Earth Orbit sorties, but as humans move further out into space, changes to this concept will have to be made. CMMD will support a close integration of logistics planning with operations planning, thus reducing the possibility of insufficient support for distant crews and, at the same time, optimizing utilization of the logistics pipeline.

Finally, CMMD will provide remote astronaut crews with decision support analysis and situational awareness tools to deal with unexpected contingencies and allow for a greater degree of autonomy in conducting their activities, especially when rapid re-planning or re-scheduling is required.

Other Potential Applications

Besides increment planners, other NASA areas (such as JSC training/simulation facilities management and Mission Control Center position scheduling) expressed interest in collaborating with the CMMD team in developing the system and adapting it to assist with their needs.

The general CMMD framework is also easily adaptable for application in other fields that require distributed, coordinated planning capabilities, and also have the need for negotiation support and contingencies investigation. Some of these include:

- *Military and civilian aviation*, where CMMD can assist with crew scheduling and equipment maintenance issues
- *Trucking industry*, where CMMD features of tightly integrating planning, scheduling, and manifesting may prove to be valuable
- *Command and control operations*, where the distributed nature of CMMD operations may help to tie together organizations with different needs/goals and provide them with the ability to coordinate through a single, detailed, always available schedule

Conclusions

The ability to generate and maintain plans and schedules in a multi-discipline, distributed environment is a subject of much research, with the need for effective applications demonstrated in many fields of human activity. This is particularly true in manned/unmanned space flight environments where planning and scheduling involves multiple distributed disciplines using a collection of disparate tools and manual techniques. Although these methods are sufficient for current NASA increment planning/scheduling operations, it will be costly and cumbersome to scale them up to meet future challenges of space exploration.

CMMD is a multi-agent planning and scheduling decision support tool being developed for future NASA space flight operations. Although targeted for the Space Exploration

Program, the initial prototypes will, as a proof-of-concept, be used to support planning and scheduling of the current International Space Station operations.

CMMD is not being designed as an autonomous planning/scheduling tool. Rather, it is a planning, scheduling, and execution decision-support system that aids planners in developing and, if needed, quickly changing their plans and schedules.

A key feature of CMMD is the integration of maintenance into the overall planning and scheduling process. Through negotiation technologies, both scheduled and unscheduled maintenance, along with its impact on the overall logistic supply-chain are considered every time a mission plan, weekly, or daily schedule is created, updated, or executed. This close integration of maintenance and logistics into every planning/scheduling operation is key to preventing unexpected mission delays and/or aborts due to maintenance or supply-chain shortcomings.

The CMMD architecture is centered on the concept of "Living Schedule" in which multiple agents – each representing a unique NASA planning discipline – interact. The Living Schedule is persistent and available to all agents on a 24x7 basis. It contains both the current and the future plans and schedules, allowing for simultaneous mission planning and execution.

Other key features of the system include generating so-called "good enough, soon enough" plans, legal schedules, localized planning/scheduling, and the ability to undertake multiple-path "what-if" analysis.

References

- [1] desJardins, M. E., Durfee, E. H., Ortiz, Jr., E. H., and Wolverton, M. J., "A Survey of Research in Distributed, Continual Planning." AI Magazine 20(4):13-22 (Special section on Distributed Continual Planning), Winter, 1999.
- [2] Ney, Z. and Looper, C., "Integrated Planning and Logistics of Exploration Initiatives for Long Term Maintenance EVA/EVR Operations", AIAA 1st Space Exploration Conference, 2004.
- [3] Ghallab, M. and Laruelle, H. "Representation and Control in IxTeT, a Temporal Planner". AIPS, 1994.
- [4] Smith, D., Frank, J., and Jonsson, A. "Bridging the Gap Between Planning and Scheduling". Knowledge Engineering Review 15(1), 2000.
- [5] Jonsson, A. and Frank, J. "A framework for dynamic constraint reasoning using procedural constraints". ECAI, 2000.
- [6] Erol, K., Hendler, J., and Nau, D., "UMCP: A Sound and Complete Procedure for Hierarchical Task-network Planning", AIPS, 1994.
- [7] Nau, D., Au, T.-C., Ilghami, O., Kuter, U., Murdock, W., Wu, D., and Yaman, F. "SHOP2: An HTN Planning System". JAIR 20:379-404, 2003.
- [8] Currie, K. and Tate, A. "O-Plan: the open planning architecture". AI 52:49-86, 1991.